HP Medical Archive Solutions
Audit Message Reference Guide

invent

March 2005 (First Edition)
Part Number 393234-001

March 2005 (First Edition)
Part Number 393234-001

DISCLAIMER

While every reasonable effort has been made to achieve technical accuracy and completeness, information in this document is subject to change without notice and does not represent a commitment on the part of Bycast Inc., or any of its subsidiaries, affiliates, licensors, or resellers. There are no warranties, express or implied, with respect to the content of this document.

Features and specifications of Bycast® products are subject to change without notice.

This manual contains information and images about Bycast Inc., its fixed content storage systems, and its other products that are protected by copyright and furnished under terms of a license agreement.

This product includes software developed by the OpenSSL Project for use in the Open SSL Toolkit. (http://www.openssl.org/)

Contents

Preface

Purpose

The Audit Management System (AMS) service stores audit messages of grid activity and events to a set of text log files. To enable you to read and analyze the audit trail, this document provides information on the structure and content of the text file log.

The objectives of this document are to:

- Describe how to access the current log file and archived logs
- Describe the text file format
- Provide a reference for common audit messages

Currency

The content is current with the AMS service software **version 4.6.0**, as included in the HP Medical Archive system **release 5.2**. To find the version number of your AMS service software:

1. Using the NMS interface, select an **AMS** service **Overview** page.

   The version number is reported in the Node Information block.

If you have an earlier version of the AMS service, contact HP Support.

Intended Audience

The content of this guide is intended for administrators responsible for producing reports of network activity and usage that require analysis of the audit messages.

You are assumed to have a sound understanding of the nature of audited activities within the HP Medical Archive system. To use the text log file, you are assumed to have access to the configured audit share on the server hosting the AMS service.

References

This document assumes familiarity with many terms related to computer operations and programming, network communications, and operating system file operations. There is wide use of acronyms. To assist you, there is a glossary at the back of this reference (page 65).

Document Structure

HP Medical Archive product guides are generally provided in printed format. They may also be available in Adobe® Acrobat® PDF (Portable Document Format).

You may print copies of the PDF editions for internal use but all copies must be treated as proprietary and confidential; *not* for general distribution.

Using this Guide

This guide is comprised of three chapters:

“Audit Message Overview”—Provides a brief overview of the audit message system and the design of the text log file.

“File and Message Format”—Defines the format of the audit log file and the format of audit messages, along with details of the common elements found in all audit messages.

“Message Reference”—Provides supporting information for all audit messages issued by the system.

Conventions

This guide adheres to conventions for terminology to avoid confusion or misunderstanding. There are also conventions for typography to enhance readability and usefulness of the text.

Terminology

There is some room for confusion between common computer network terminology for “server” and “node” as they are used in HP Medical Archive products and documents.

A server is usually thought of as a piece of computing hardware that provides data services to requesting network clients; a resource providing network, computational, and storage services. Within the context of the HP Medical Archive, a server is an entity hosting one or more grid services.

Nodes in a network are usually defined as an independent entity with a unique network identity, running on a resource. In this text, the use of the phrase “grid node” refers to an addressable entity on the grid that provides and uses functional services within the grid to perform one or more tasks. Each grid node has a unique “node ID”. These include: ADC, CMS and LDR. In the HP Medical Archive User Guide and other user documents these are referred to as “services”.

In contrast, the HP Medical Archive packages the grid service modules into “nodes”. Some node packages are required, others are optional. When used in this context the term appears in uppercase; as in “ControlNODE”, which usually incorporates the ADC, CMS and SSM services on one server.

Numerics

Numeric values are presented in decimal unless noted otherwise.

Hexadecimal values in the narrative are noted using the prefix “0x”; for example: 0x3B. Where sample messages include data as a string of hexadecimal characters, the prefix only appears if it is included within the message.

Fonts

To assist you in easily picking out the elements of importance, changes from the standard font are used:

- Items upon which you act are shown in bold. These include:
    - Sequences of selections from the navigation tree, tabs, and page options, such as: **LDR ▶ Configuration ▶ Notifications**.
    - Buttons or keys to click or press, such as **Apply** or **<Tab>**.
    - Radio buttons or check buttons to enable or disable, such as **Save configuration as default**.

- Field prompts, names of windows and dialogs, messages, and other literal text in the interface is shown in sans-serif such as the LDR State pull down menu, or the Sign In... window.
- Items within the narrative that require emphasis appear in *italics*.
- Coding samples or interactions with a command terminal are shown in the fixed space font: `<?xml version=1.0 ?>`

  Any italicized portion indicates variable data you provide to meet your needs.

Keyboard keys that use words or standard abbreviations are shown within angle brackets, such as **<Ctrl>** for the control key, **<Tab>**, **<space>**, and **<Enter>**.

Contacts

For general product and company information, refer to the HP web site at:

www.hp.com

If you cannot find the information you need in this document, there are several other resources you can use to get more detailed information:

- The HP website (http://www.hp.com)
- Your nearest HP authorized reseller (for the locations and telephone numbers of these resellers, refer to the HP website)
- HP technical support:
  - In North America, call 1-800-652-6672
  - For other regions, refer to the HP website.

1

Audit Message Overview

Chapter Contents

Overview of Auditing

As services in the grid perform various activities and process events, audit messages are generated to retain a record of grid activity. These messages are processed by the Audit Management System (AMS) service and stored in the form of text log files. This document provides information on the structure and content of the text log files to enable you to read and analyze the audit trail of grid activity.

Audit Message Flow

Audit messages are generated internally by each grid service. All system services generate audit messages during normal system operation. These messages are sent to the connected AMS services for processing and storage.

Some grid services can be designated as audit message relay services. They act as collection points to reduce the need for every service to send its audit messages to all connected AMS services. Notice in Figure 1 that each relay service must send messages to all AMS destinations, whereas services can send messages to just one relay service.

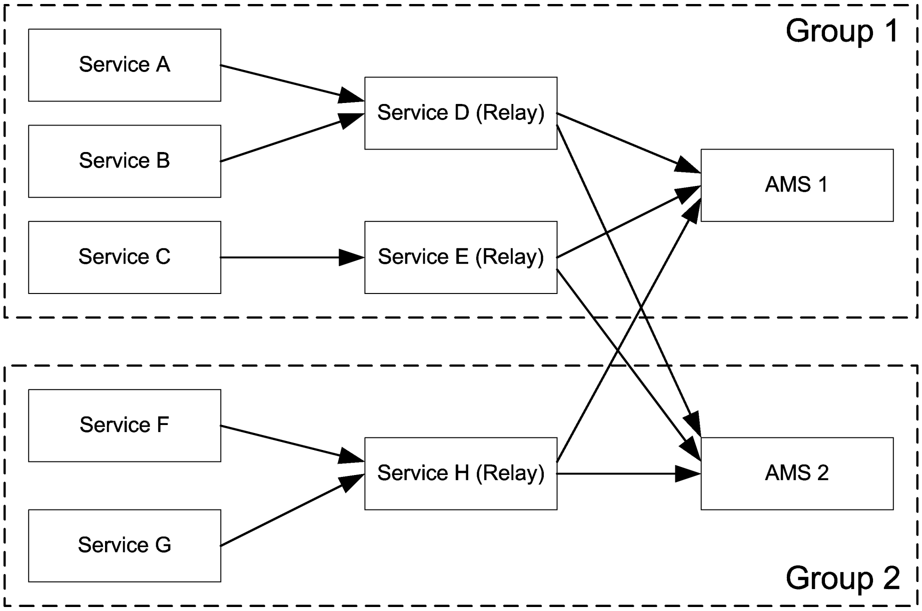

Figure 1: Audit Message Flow

Relay services are designated at the time the grid topology is configured. Any grid service (LDR, ADC, CMS, and so on) can be designated to act as an audit message relay.

Message Retention

Once an audit message is generated, it is stored on the local server of the originating service until it has been committed to all connected AMS servers, or a designated audit relay service. The relays in turn store the message until it is committed at all AMS services. This process includes a confirmation (positive acknowledgment) to ensure no messages are lost.

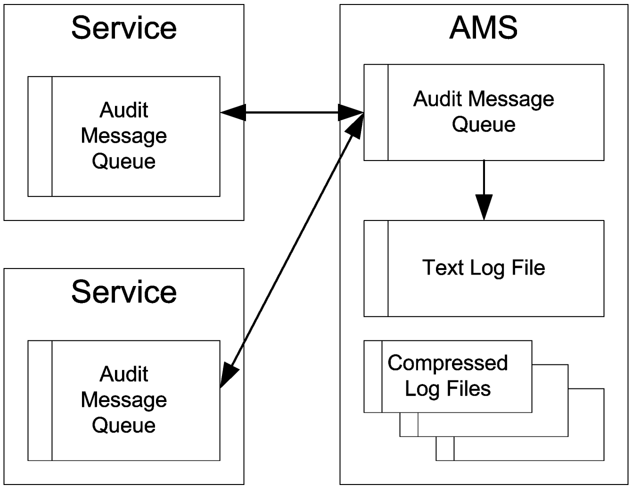

Figure 2: Audit Message Retention

Messages arrive at the AMS and are stored in a queue pending confirmed write to the text log file. Confirmation of the arrival of messages is sent to the originating service (or audit relay) to permit the originator to delete its copy of the message.

Only after a message has been committed to storage at the AMS can it be removed from the queue. This local message buffer at the AMS has an alarm (AMQS) associated with it, in the event the backlog becomes unusually large. At times of peak activity, the rate at which audit messages are arriving may be faster than they can be committed to storage, causing a temporary backlog that will clear itself when grid activity declines.

When the text log file on the Admin Node reaches a predefined size, it is automatically converted to a compressed format and a new text log file is started. Over very long periods of time, this can result in consumption of the available storage on the server hosting the AMS service. Based on the requirements of your enterprise, either archive the older compressed files to some other media (such as DVD-R, or into the grid itself), or they will be automatically deleted.

Audit Log File Access

Access to the text log file at the AMS requires you to have an account and password to access the audit share on the server hosting the AMS service.

The active log file and any compressed log files are available through your configured audit share directory.

The active audit log file is named:

```
audit.log
```

Archived log files are named using the convention:

```
YYYY-MM-DD.txt.gz
```

where the file name includes a date and time stamp (in UTC) when the file was archived.

To access an archived audit log file:

1. Make a local copy of the file to work with.
2. Decompress the file. This process requires a decompression utility. We recommend "7-Zip", which is a free download from:
   http://www.7-zip.org/

Access log files as simple text files.

The next chapter provides details of the file's internal structure and the syntax of audit messages.

File and Message Format

2

Chapter Contents

Audit Log File Format

The audit log contains individual audit messages in the following format:

1. Date and time stamp (local time) the message was processed at the AMS, followed by the server host name and the string “AMS:”.
2. The message itself, enclosed within square brackets “[]”. The message structure is discussed in the next section on page 6.

The following is the beginning of a sample log file. Messages are wrapped within the boundaries shown, ending after the ASQN attribute and double closing brackets “]]”. The <CR><LF> characters at the end of each message are not shown.

```
Feb 12 02:37:34 an1-a-1 AMS:
[AUDT[RSLT(FC32):'DSDN'][AVER(UI32):3][ATYP(FC32):'SYSU'][ATIM(UI64):11081758444743
62][ATID(UI64):9384121014334693630][ANID(UI32):15010119][AMID(FC32):'ARNI'][ASQN(UI
64):0]]
Feb 12 02:37:34 an1-a-1 AMS:
[AUDT[SEID(FC32):'RCON'][CNDR(FC32):'OUTB'][SVIP(UI32):1501][DAIP(IP32):14.1.1.13][
SAIP(IP32):14.1.1.19][CNID(UI64):1716307103][RSLT(FC32):'CRFU'][AVER(UI32):3][ATYP(
FC32):'ETCF'][ATIM(UI64):1108175844660669][ATID(UI64):5503182624165676149][ANID(UI3
2):15010119][AMID(FC32):'RCON'][ASQN(UI64):1]]
Feb 12 02:37:34 an1-a-1 AMS:
[AUDT[SEID(FC32):'RCON'][CNDR(FC32):'OUTB'][SVIP(UI32):1501][DAIP(IP32):14.1.1.15][
SAIP(IP32):14.1.1.19][CNID(UI64):2329159112][RSLT(FC32):'CRFU'][AVER(UI32):3][ATYP(
FC32):'ETCF'][ATIM(UI64):1108175854682710][ATID(UI64):7756750787035320318][ANID(UI3
2):15010119][AMID(FC32):'RCON'][ASQN(UI64):2]]
```

Audit Message Format

Audit messages exchanged within the grid include some standard information common to all messages, and specific content for the event or activity being reported.

Each audit message is logged as a string composed of attribute elements that are:

- Enclosed in square brackets “[ ]”
- Introduced by the string “AUDT”, indicating an audit message

- Do not have delimiters (no commas or spaces) between attributes
- Terminated by a carriage return and line feed (<CR><LF>)

Each element includes: an attribute code, data type, and value. It takes the format:

```
[ATTR(type):value][ATTR(type):value]...
[ATTR(type):value]<CR><LF>
```

Where:

- **`ATTR`** is a four-character code for the attribute being reported. See Chapter 3, starting on page 11 for a directory of message attributes and their meaning.
- **`type`** is a four-character identifier of the programming data type of the value, such as: UI64, FC32, and so on. See "Data Types" on the next page. The type is enclosed in brackets "( )".
- **`value`** is the content of the attribute, typically a numeric or text value (text values are enclosed in single quotes). Values always follow a colon ":".

The number of attribute elements in the message depends on the event type of the message.

Sample Audit Message

```
[HSID(UI64):811028912][DIDR(CSTR):'INBO'][HSCR(FC32):'SUCS']
[AVER(UI32):2][ATYP(FC32):'HTSC'][ATIM(UI64):1099615457414746]
[ATID(UI64):11174705928149966150][ANID(UI32):12130010]
[AMID(FC32):'HTSM'][ASQN(UI32):49984]
```

Data Types

The data types encountered in the audit messages are:

Table 1: Data Types

Type	Description
UI32	Unsigned long integer (32 bits); it can store the numbers 0–4,294,967,295.
UI64	Unsigned double long integer (64 bits); it can store the numbers 0–18,446,744,073,709,551,615.
FC32	Four Character Constant; a 32-bit unsigned integer value represented as four ASCII characters such as: "ABCD".

Table 1: Data Types (cont.)

Type	Description
IP32	IP Address; a 32-bit IP address representation.
CSTR	C String; a variable length array of characters.

Event-Specific Data

Following the opening "[AUDT" container that identifies the message itself, is a series of items specific to each event or action. Chapter 3, "Message Reference" on page 11 lists attributes commonly used for tracing grid activity.

Common Elements

After the event-specific information is a set of elements common to all audit messages:

Table 2: Common Elements of Audit Messages

Code	Type	Description
AVER	UI32	Version—The version of the audit message. As the HP Medical Archive software evolves, new versions of services may incorporate new features in audit reporting. This field enables backward compatibility in the AMS to process messages from older versions of services.
ATYP	FC32	Event Type—A four-character identifier of the event being logged. This governs the "payload" content of the message—the attributes included.
ATIM	UI64	Timestamp—The time the event was generated that triggered the audit message, measured in microseconds since the operating system epoch (00:00:00 UTC on 1 January, 1970). Note that most available tools for converting the timestamp to local date and time are based on milliseconds. Rounding or truncation of the database timestamp may be required.
ATID	UI64	Trace ID—An identifier that is shared by the set of messages that were triggered by a single event.

Table 2: Common Elements of Audit Messages (cont.)

Code	Type	Description
ANID	UI32	Node ID—The grid node ID assigned to the service that generated the message. Each service is allocated a unique identifier at the time the HP Medical Archive is configured and installed. This ID cannot be changed.
AMID	FC32	Module ID—A four-character identifier of the module ID that generated the message. This indicates the code segment within which the audit message was generated.
ASQN	UI64	Sequence Count—A counter that is incremented for each generated audit message on the grid node (ANID). This counter is reset to zero at service restart. It can be used for consistency checks to ensure that no audit messages have been lost.

Message Reference

3

A comprehensive listing of generated audit messages.

Chapter Contents

Introduction

This chapter provides detailed descriptions of the attributes reported in all audit messages issued by the system.

Messages are listed alphabetically to facilitate referencing the content for a specific message of interest. To reference related messages for a given class of activity, use the tables in the subsections below.

System Audit Messages

This group of messages are for events related to:

- The auditing system itself
- Grid node states
- Grid-wide task activity (Grid Tasks)
- Service backup operations
- File System Gateway (FSG) replications

Table 3: System Audit Messages

Code	Description	Page
ETCA	TCP/IP Connection Establish—An incoming or outgoing TCP/IP connection was successfully established.	37
ETCC	TCP/IP Connection Close—An established connection has been closed by either side of the connection (normally or abnormally).	38
ETCF	TCP/IP Connection Fail—An outgoing connection attempt failed at the lowest level, due to communication problems.	38
SADD	Security Audit Disable—Audit message logging has been turned off.	56
SADE	Security Audit Enable—Audit message logging has been turned on.	57
ETAF	Security Authentication Failed—A connection attempt using Transport Layer Security (TLS) has failed.	36
SYSU	Node Start—An HP Medical Archive grid service started; the nature of the previous shutdown is indicated in the message.	60
SYSD	Node Stop—An HP Medical Archive grid service has been gracefully stopped.	57

Table 3: System Audit Messages (cont.)

Code	Description	Page
TSTC	Grid Task State Change—A grid task has been added, started, paused, canceled, or completed.	63
TSGC	Grid Task Stage Change—The stage of a grid task has changed.	62
TACB	Grid Task Action Begin—A grid task action has begun.	61
TACE	Grid Task Action End—A grid task action has completed.	61
BKSB	Backup Store Begin—A service has begun a backup operation.	17
BKSE	Backup Store End—A service has completed a backup operation.	18
RPSB	Replication Session Begin—A service has begun a replication operation to a secondary service.	55
RPSE	Replication Session End—A service has completed a replication operation to a secondary service.	55

Object Audit Messages

Object audit messages represent events related to the storage and management of objects within the grid. These include:

- Object storage/retrieval
- Node-to-node transfer
- Verification

Table 4: Object Audit Messages

Code	Description	Page
CBSB	Object Send Begin—The source entity initiated a node-to-node data transfer operation on a single piece of content.	21
CBSE	Object Send End—The source entity completed a node-to-node data transfer operation.	22
CBRB	Object Receive Begin—The destination entity initiated a node-to-node data transfer operation on a single piece of content.	19
CBRE	Object Receive End—The destination entity completed a node-to-node data transfer operation.	20
SCMT	Object Store Commit—A content block was completely stored and verified, and can now be requested.	57

Table 4: Object Audit Messages

Code	Description	Page
SREM	Object Store Remove—A content block was deleted from a node, and can no longer be requested directly.	58
SVRF	Object Store Verify Fail—A content block failed verification checks.	59
SVRU	Object Store Verify Unknown—Unexpected file(s) detected in the object store.	59

HTTP Protocol Audit Messages

HTTP Protocol audit messages represent events related to interactions with internal and external system components using the HTTP protocol. These include:

- Session establishment/breakdown
- Object storage
- Retrieval
- Query

Table 5: HTTP Protocol Audit Messages

Code	Description	Page
HTSE	HTTP Session Establish—A remote host successfully established an HTTP session to the node.	54
HTSC	HTTP Session Close—An HTTP client closed a previously-established HTTP session.	53
HHEA	HTTP HEAD Transaction—Information about a piece of content was requested by an HTTP client.	48
HGES	HTTP GET Transaction Start—A request for a GET transaction to transfer content to an HTTP client was initiated.	47
HGEE	HTTP GET Transaction End—A GET transaction to transfer content to an HTTP client completed.	46
HPUS	HTTP PUT Transaction Start—A PUT transaction to transfer content from an HTTP client was initiated.	53
HPUE	HTTP PUT Transaction End—A PUT transaction to transfer content from an HTTP client completed.	52

Table 5: HTTP Protocol Audit Messages

Code	Description	Page
HPOS	HTTP POST Transaction Start—An HTTP client initiated a query for stored content.	51
HPOE	HTTP POST Transaction End—An HTTP client completed a query for stored content.	50
HDEL	HTTP DELETE Transaction—Logs the result of a request to delete content.	45
HOPT	HTTP OPTIONS Transaction—Logs the result of a request for information about the transactions that can be performed on content.	49
HCPS	HTTP PUT C–STORE Start—A PUT transaction to transfer content between hosts was initiated.	45
HCPE	HTTP PUT C–STORE End—A PUT transaction to transfer content between hosts completed.	44

DICOM Audit Messages

This set of messages log activity related to interactions with external systems using the DICOM protocol. These include:

- Association establishment
- C–STORE
- C–FIND
- C–MOVE
- N–ACTION (storage commitment)

Table 6: DICOM Audit Messages

Code	Description	Page
DASE	DICOM Association Establish—A successful inbound or outbound DICOM association was established with a remote host.	24
DASC	DICOM Association Close—An established DICOM association with a remote host closed.	24
DASF	DICOM Association Fail—An association attempt failed (remote host cannot process the DICOM protocol, or the request was rejected).	25
DCPS	DICOM C–STORE Start—A transfer of content between hosts over a DICOM association has started.	34

Table 6: DICOM Audit Messages

Code	Description	Page
DCPE	DICOM C–STORE End—A transfer of content between hosts over a DICOM association has completed.	32
DCSF	DICOM C–STORE Fail—A transfer of content between hosts over a DICOM association has failed.	34
DCFS	DICOM C–FIND Start—A remote DICOM host initiated a query for DICOM-related content.	27
DCFE	DICOM C–FIND End—A remote DICOM host completed a query for DICOM-related content.	26
DCGS	DICOM C–GET Start—A remote DICOM host initiated a query/ retrieve for DICOM-related content.	29
DCGE	DICOM C–GET End—A remote DICOM host completed a query/ retrieve for DICOM-related content.	28
DCMS	DICOM C–MOVE Start—A remote DICOM host initiated a transfer of DICOM instances to a remote Application Entity.	31
DCME	DICOM C–MOVE End—A remote DICOM host completed a transfer of DICOM instances to a remote Application Entity.	29
DCMT	DICOM Storage Commitment—A remote DICOM host initiated an operation to check if content was previously stored.	32
CDAD	DICOM Study Add—A new study (not previously recorded by the CMS) or a new instance (image) to a known study has been added.	23

File System Gateway Audit Messages

This set of messages log activity related to interactions with external systems via the File System Gateway (FSG) interface to the grid.

Table 7: File System Gateway Audit Messages

Code	Description	Page
FCRE	File Create—Logs the addition of new files (not directories) to the FSG.	40
FDEL	File Delete—Logs deletion of a file from the FSG directory tree (not from the grid).	40
FRNM	File Rename—Logs changes to the name or path of an existing file.	41

Table 7: File System Gateway Audit Messages

Code	Description	Page
FMFY	File Modify—Logs changes to the content of an existing file.	41
FSTG	File Store to Grid—Logs the storage of content from the FSG local cache to the grid.	42
FSWO	File Swap Out—Logs the deletion of a file from the FSG local cache (but not from the directory tree or grid).	43
FSWI	File Swap In—Logs the retrieval of a file from the grid to the FSG local cache.	42

As content is added to the grid via the FSG, the content is first stored locally in a cache on the FSG server. The FSG manages ingesting the content to the grid. The content in the cache can be purged if space is needed for new content, either inbound or outbound. As the cache content is changed, additional audit messages are logged.

Any changes made to the name or content of a file previously entered in the FSG are also logged, as are file deletions from the FSG. Note that deletions to the FSG result in removal of the entry from the FSG directory tree; however the file content is retained in the grid and can be directly accessed via the assigned content block ID.

Audit Message Reference

BKSB—Backup Store Begin

When a service begins a backup operation—storing private structured data to the grid—this message is generated.

Table 8: BKSB—Backup Store Begin Fields

Code	Field	Description
BKSI	Backup Session ID	The unique identifier of the backup session that is being started.
BKOI	Backup Source Entity	The type of entity that is performing the backup; typically one of: BFSG, BCMS, or BNMS.

Table 8: BKSB—Backup Store Begin Fields (cont.)

Code	Field	Description
BKEE	Entries to Backup	The number of entries (objects) the entity expects to include in this backup session. If the value is unknown, this field is set to zero (0).
RSLT	Backup Initiation Status	This field indicates status at the time the backup store was initiated: SUCS—the backup store started successfully.

This message marks the time of a backup session. It allows you to match the message with a corresponding BKSE end message to determine that backups are happening as planned and whether they are successful.

BKSE—Backup Store End

When a service completes a backup operation, this message is generated.

Table 9: BKSE—Backup Store End Fields

Code	Field	Description
BKSI	Backup Session ID	The unique identifier of the backup session that has been completed.
BKOI	Backup Source Entity	The type of entity that performed the backup; typically one of: BFSG, BCMS, or BNMS.
BKEA	Entries Backed Up	The actual number of entries (objects) that were included in this backup session. You can compare this to BKEE in the BKSB message.
UUID	Backup UUID	The Universal Unique IDentifier assigned to the backup by the grid. If the backup session fails or is aborted, this value is the NULL UUID.
RSLT	Backup Result	The completion status of the backup session: SUCS—The backup completed successfully. ABRT—The backup was aborted. FAIL—The backup failed before completion. STFL—The backup data could not be stored in the grid.

Matching this message with the corresponding BKSB message can indicate the time it took to perform the backup. This message indicates whether the backup was successful and the UUID of the backup data within the grid, should a restoration be needed.

CBRB—Object Receive Begin

During normal system operations, content blocks are continuously transferred between different nodes as data is accessed, replicated and retained. When transfer of a content block from one node to another is initiated, this message is issued by the destination entity.

Table 10: CBRB—Object Receive Begin Fields

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the node-to-node content block transfer.
CBID	Content Block Identifier	The unique identifier of the content block being transferred.
CTDR	Transfer Direction	Indicates if the CBID transfer was push-initiated or pull-initiated: PUSH—the transfer operation was requested by the sending entity. PULL—the transfer operation was requested by the receiving entity.
CTSR	Source Entity	The node ID of the source (sender) of the CBID transfer.
CTDS	Destination Entity	The node ID of the destination (receiver) of the CBID transfer.
CTSS	Start Sequence Count	Indicates the first sequence count requested. If successful, the transfer begins from this sequence count.
CTES	Expected End Sequence Count	Indicates the last sequence count requested. If successful, the transfer is considered complete when this sequence count has been received.
RSLT	Transfer Start Status	Status at the time the transfer was started: SUCS—transfer started successfully.

This audit message means a node-to-node data transfer operation was initiated on a single piece of content, as identified by its Content Block Identifier. The operation requests data from “Start Sequence Count” to “Expected End Sequence Count”. Sending and receiving nodes are identified by their node IDs. This information can be used to track system data flow, and when combined with storage audit messages, to verify replica counts.

CBRE—Object Receive End

When transfer of a content block from one node to another is completed, this message is issued by the destination entity.

Table 11: CBRE—Object Receive End Fields

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the node-to-node content block transfer.
CBID	Content Block Identifier	The unique identifier of the content block being transferred.
CTDR	Transfer Direction	Indicates if the CBID transfer was push-initiated or pull-initiated: PUSH—the transfer operation was requested by the sending entity. PULL—the transfer operation was requested by the receiving entity.
CTSR	Source Entity	The node ID of the source (sender) of the CBID transfer.
CTDS	Destination Entity	The node ID of the destination (receiver) of the CBID transfer.
CTSS	Start Sequence Count	Indicates the sequence count on which the transfer started.
CTAS	Actual End Sequence Count	Indicates the last sequence count successfully transferred. If the Actual End Sequence Count is the same as the Start Sequence Count, and the Transfer Result was not successful, no data was exchanged.
RSLT	Transfer Result	The result of the transfer operation (from the perspective of the sending entity): SUCS—transfer successfully completed; all requested sequence counts were sent. CONL—connection lost during transfer CTMO—connection timed-out during establishment UNRE—destination node ID unreachable CRPT—transfer ended due to reception of corrupt or invalid data (may indicate tampering)

This audit message means a node-to-node data transfer operation was completed. If the Transfer Result was successful, the operation transferred data from “Start Sequence Count” to “Actual End Sequence Count”. Sending and receiving nodes are identified by their node IDs. This information can be used to track system data flow and to locate, tabulate, and analyze errors. When combined with storage audit messages, it can also be used to verify replica counts.

CBSB—Object Send Begin

During normal system operations, content blocks are continuously transferred between different nodes as data is accessed, replicated and retained. When transfer of a content block from one node to another is initiated, this message is issued by the source entity.

Table 12: CBSB—Object Send Begin Fields

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the node-to-node content block transfer.
CBID	Content Block Identifier	The unique identifier of the content block being transferred.
CTDR	Transfer Direction	Indicates if the CBID transfer was push-initiated or pull-initiated: PUSH—the transfer operation was requested by the sending entity. PULL—the transfer operation was requested by the receiving entity.
CTSR	Source Entity	The node ID of the source (sender) of the CBID transfer.
CTDS	Destination Entity	The node ID of the destination (receiver) of the CBID transfer.
CTSS	Start Sequence Count	Indicates the first sequence count requested. If successful, the transfer begins from this sequence count.
CTES	Expected End Sequence Count	Indicates the last sequence count requested. If successful, the transfer is considered complete when this sequence count has been received.
RSLT	Transfer Start Status	Status at the time the transfer was started: SUCS—transfer started successfully.

This audit message means a node-to-node data transfer operation was initiated on a single piece of content, as identified by its Content Block Identifier. The operation requests data from “Start Sequence Count” to “Expected End Sequence Count”. Sending and receiving nodes are identified by their node IDs. This information can be used to track system data flow, and when combined with storage audit messages, to verify replica counts.

CBSE—Object Send End

When transfer of a content block from one node to another is completed, this message is issued by the source entity.

Table 13: CBSE—Object Send End Fields

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the node-to-node content block transfer.
CBID	Content Block Identifier	The unique identifier of the content block being transferred.
CTDR	Transfer Direction	Indicates if the CBID transfer was push-initiated or pull-initiated: PUSH—the transfer operation was requested by the sending entity. PULL—the transfer operation was requested by the receiving entity.
CTSR	Source Entity	The node ID of the source (sender) of the CBID transfer.
CTDS	Destination Entity	The node ID of the destination (receiver) of the CBID transfer.
CTSS	Start Sequence Count	Indicates the sequence count on which the transfer started.
CTAS	Actual End Sequence Count	Indicates the last sequence count successfully transferred. If the Actual End Sequence Count is the same as the Start Sequence Count, and the Transfer Result was not successful, no data was exchanged.

Table 13: CBSE—Object Send End Fields (cont.)

Code	Field	Description
RSLT	Transfer Result	The result of the transfer operation (from the perspective of the sending entity): SUCS—transfer successfully completed; all requested sequence counts were sent. CONL—connection lost during transfer CTMO—connection timed-out during establishment UNRE—destination node ID unreachable CRPT—transfer ended due to reception of corrupt or invalid data (may indicate tampering)

This audit message means a node-to-node data transfer operation was completed. If the Transfer Result was successful, the operation transferred data from “Start Sequence Count” to “Actual End Sequence Count”. Sending and receiving nodes are identified by their node IDs. This information can be used to track system data flow and to locate, tabulate, and analyze errors. When combined with storage audit messages, it can also be used to verify replica counts.

CDAD—DICOM Study Add

When a new DICOM study ID is ingested, or when new images are added to an existing study, this logs the addition.

Table 14: CDAD—DICOM Study Add Fields

Code	Field	Description
STDY	Study GUID	The unique DICOM study identifier.
SIMC	Number of Images	The number of instances (images) in the study.

This audit message appears for each new study instance (image) that is added to the grid. As a new study appears, the message indicates the new study is now known to the grid. As images are added to the study, the message appears with the new count of images.

DASC—DICOM Association Close

When an established DICOM association with a remote host is closed, this message is issued.

Table 15: DASC—DICOM Association Close Fields

Code	Field	Description
ASID	Association Identifier	The unique identifier assigned to the DICOM association.
RSLT	Closing State	Indicates how the association closed: SUCS—closed normally without errors TOUT—timed-out by the node due to inactivity ERRC—lost connection ABRT—aborted GERR—general data processing error

This audit message means the DICOM association specified by the Association Identifier is no longer established. The DASC message always corresponds with a previous DASE (Association Establish) message. DASC should be monitored to determine if there are excessive problems during attempts to establish an association. Problems could indicate communications or interoperability issues related to DICOM implementation.

DASE—DICOM Association Establish

When a DICOM association is established between a node and a host, this message is issued.

Table 16: DASE—DICOM Association Establish Fields

Code	Field	Description
CNID	Connection Identifier	The unique identifier assigned to the connection over which the DICOM association was established.
ASID	Association Identifier	The unique identifier assigned to the DICOM association.

Table 16: DASE—DICOM Association Establish Fields (cont.)

Code	Field	Description
DIDR	Association Direction	Indicates whether the association was opened by the grid node or by a remote host: INBO—initiated by a remote host connecting to the grid node OUTB—initiated by the grid node connecting to a remote host
RMAE	External Application Entity	The Application Entity Title of the remote device.
GRAE	Grid Application Entity	The Application Entity Title of the grid.

This audit message means a successful inbound or outbound DICOM association was established with a remote host. It can be used to track hosts communicating with the system via DICOM.

The Grid Application Entity field allows identification of related configuration and coerce tag profiles, if applicable.

DASF—DICOM Association Fail

When an attempt by a DICOM service to establish an association fails, this message is issued. This can occur if the remote host cannot process the DICOM protocol, or when either side of the communication rejects the association request.

Table 17: DASF—DICOM Association Fail Fields

Code	Field	Description
CNID	Connection Identifier	The unique identifier assigned to the connection over which the DICOM association was established.
ASID	Association Identifier	The unique identifier assigned to the DICOM association.
DIDR	Association Direction	Indicates whether the association was opened by the grid node or by a remote host: INBO—initiated by a remote host connecting to the grid node OUTB—initiated by the grid node connecting to a remote host

Table 17: DASF—DICOM Association Fail Fields (cont.)

Code	Field	Description
RMAE	External Application Entity	The Application Entity Title of the remote device (if unknown, this field contains a null string).
GRAE	Grid Application Entity	The Application Entity Title of the grid.
RSLT	Failure Code	Reason for the failure: ERRC—connection closed by remote host before an association could be established TOUT—timeout period expired REJT—association rejected PERM—calling AE Title denied permission to connect CONF—unexpected remote AE Title COMP—suitable presentation context could not be negotiated GERR—unknown data received from remote host

This audit message should be monitored to determine if there are repetitive or excessive problems during attempts to establish an association. Problems could indicate communications or interoperability issues related to DICOM implementation, or incorrectly configured external DICOM devices.

The Grid Application Entity field allows identification of related configuration and coerce tag profiles, if applicable.

DCFE—DICOM C–FIND End

When a DICOM association completes a C–FIND operation to query available content, this message is issued.

Table 18: DCFE—DICOM C–FIND End Fields

Code	Field	Description
ASID	Association Identifier	The unique identifier assigned to the DICOM association.
DIDR	C–FIND Direction	Indicates whether the C–FIND was initiated by the grid node or by a remote host: INBO—initiated by a remote host

Table 18: DCFE—DICOM C–FIND End Fields (cont.)

Code	Field	Description
ROOT	DICOM Query Root	The query root specified in the C–FIND.
LEVL	DICOM Query Level	The query level specified in the C–FIND.
RSFD	Results Found	The number of DICOM objects found matching the query.
RSLT	Result Code	The result of the C–FIND operation: SUCS—successful CANC—cancelled by the Service Class User GERR—general error processing the C–FIND command

This audit message means a remote DICOM host initiated and completed a query for DICOM-related content. It can be monitored to determine the content being queried. The "Result Code" field can be used to determine when errors occur.

The time interval between the C–FIND Start and C–FIND End audit messages tells you how long the related C–FIND operations are taking to complete.

DCFS—DICOM C–FIND Start

When a DICOM association initiates a C–FIND operation to query available content, this message is issued.

Table 19: DCFS—DICOM C–FIND Start Fields

Code	Field	Description
ASID	Association Identifier	The unique identifier assigned to the DICOM association.
DIDR	C–FIND Direction	Indicates whether the C–FIND was initiated by the grid node or by a remote host: INBO—initiated by a remote host
ROOT	DICOM Query Root	The query root specified in the C–FIND.
LEVL	DICOM Query Level	The query level specified in the C–FIND.

This audit message means a remote DICOM host initiated a query for DICOM-related content. It can be monitored to determine the content being queried.

The time interval between the C–FIND Start and C–FIND End audit messages tells you how long the related C–FIND operations are taking to complete.

DCGE—DICOM C–GET End

When a DICOM association completes a C–GET operation to query and retrieve found content, this message is issued.

Table 20: DCGE—DICOM C–GET End Fields

Code	Field	Description
ASID	Association Identifier	The unique identifier assigned to the DICOM association.
DIDR	C–GET Direction	Indicates whether the C–GET was initiated by the grid node or by a remote host: INBO—initiated by a remote host
ROOT	DICOM Query Root	The query root specified in the C–GET.
LEVL	DICOM Query Level	The query level specified in the C–GET.
RSFD	Results Found	The number of DICOM objects retrieved matching the query.
RSLT	Result Code	The result of the C–GET operation: SUCS—successful CANC—cancelled by the Service Class User GERR—general error processing the C–GET command

This audit message means a remote DICOM host initiated and completed a query/retrieve for DICOM-related content. It can be monitored to determine the content being retrieved. The “Result Code” field can be used to determine when errors occur.

The time interval between the C–GET Start and C–GET End audit messages tells you how long the related C–GET operations are taking to complete.

DCGS—DICOM C–GET Start

When a DICOM association initiates a C–GET operation to query and retrieve DICOM content, this message is issued.

Table 21: DCGS—DICOM C–GET Start Fields

Code	Field	Description
ASID	Association Identifier	The unique identifier assigned to the DICOM association.
DIDR	C–GET Direction	Indicates whether the C–GET was initiated by the grid node or by a remote host: INBO—initiated by a remote host
ROOT	DICOM Query Root	The query root specified in the C–GET.
LEVL	DICOM Query Level	The query level specified in the C–GET.

This audit message means a remote DICOM host initiated a query/ retrieve for DICOM-related content. It can be monitored to determine the content being retrieved. The C–STORE audit messages on the same association (i.e. C–STORE audit messages with the same Association Identifier) between the C–GET Start and C–GET End messages correspond to the retrieved objects associated with the initial query.

The time interval between the C–GET Start and C–GET End audit messages tells you how long the related C–GET operations are taking to complete.

DCME—DICOM C–MOVE End

When a DICOM association completes a C–MOVE operation to query and retrieve found content over a second association, this message is issued.

Table 22: DCME—DICOM C–MOVE End Fields

Code	Field	Description
ASID	Association Identifier	The unique identifier assigned to the DICOM association.

Table 22: DCME—DICOM C–MOVE End Fields (cont.)

Code	Field	Description
DIDR	C–MOVE Direction	Indicates whether the C–MOVE was initiated by the grid node or by a remote host: INBO—initiated by a remote host
ROOT	DICOM Query Root	The query root specified in the C–MOVE.
LEVL	DICOM Query Level	The query level specified in the C–MOVE.
SAET	Source Application Entity (AE) Title	The source AE-Title for the C–MOVE operation.
DAET	Destination AE-Title	The destination AE-Title for the C–MOVE operation.
RSFD	Results Found	The number of DICOM objects retrieved matching the query.
RSLT	Result Code	The result of the C–MOVE operation: SUCS—successful CANC—cancelled by the Service Class User GERR—general error processing the C–MOVE command

This audit message means a remote DICOM host initiated and completed a a C–MOVE operation to transfer DICOM content. It can be monitored to determine the content being queried/transferred. The "Result Code" field can be used to determine when errors occur.

The time interval between the C–MOVE Start and C–MOVE End audit messages tells you how long the related C–MOVE operations are taking to complete.

DCMS—DICOM C–MOVE Start

When a DICOM association initiates a C–MOVE operation to query and transfer DICOM content over a second association, this message is issued.

Table 23: DCMS—DICOM C–MOVE Start Fields

Code	Field	Description
ASID	Association Identifier	The unique identifier assigned to the DICOM association.
DIDR	C–MOVE Direction	Indicates whether the C–MOVE was initiated by the grid node or by a remote host: INBO—initiated by a remote host
ROOT	DICOM Query Root	The query root specified in the C–MOVE.
LEVL	DICOM Query Level	The query level specified in the C–MOVE.
SAET	Source Application Entity (AE) Title	The source AE-Title for the C–MOVE operation.
DAET	Destination AE-Title	The destination AE-Title for the C–MOVE operation.

This audit message means a remote DICOM host initiated a C–MOVE operation to transfer DICOM instances to a remote Application Entity. It can be monitored to determine the content being retrieved. The C–STORE audit messages on the same association (i.e. C–STORE audit messages with the same Association Identifier) resulting from the C–MOVE correspond to the retrieved objects.

The time interval between the C–MOVE Start and C–MOVE End audit messages tells you how long the related C–MOVE operations are taking to complete.

DCMT—DICOM Storage Commitment

When a DICOM association initiates a Storage Commitment operation to determine if content has been successfully received and stored, this message is issued.

Table 24: DCMT—DICOM Storage Commitment Fields

Code	Field	Description
ASID	Association Identifier	The unique identifier assigned to the DICOM association.
ISTR	Items Requested	The number of items requested for storage verification.
ISTS	Items Stored	The number of items requested for verification which have been successfully stored.
ISTN	Items not Stored	The number of items requested for verification which have *not* been successfully stored.
RSLT	Result Code	Result of the Storage Commitment operation: SUCS—successful GERR—an error occurred during Storage Commitment processing

This audit message means a remote DICOM host initiated a Storage Commitment operation to check whether content has been previously stored. It can be used to discover situations where a discrepancy exists between content storage requests and what was in fact successfully stored.

DCPE—DICOM C–STORE End

When a DICOM association completes a C–STORE operation to transfer content from one host to another, this message is issued.

Table 25: DCPE—DICOM C–STORE End Fields

Code	Field	Description
ASID	Association Identifier	The unique identifier assigned to the DICOM association.
DIDR	C–STORE Direction	Indicates whether the C–STORE was initiated by the grid node or by a remote host: INBO—initiated by a remote host OUTB—initiated by the node

Table 25: DCPE—DICOM C–STORE End Fields (cont.)

Code	Field	Description
STUG	Study Instance UID	The Study Identifier of the data being transferred.
SERG	Series Instance UID	The Series Identifier of the data being transferred.
IMGG	SOP Instance UID	The Image Identifier of the data being transferred.
STCL	SOP Class	The SOP Class of the instance.
STTX	Transfer Syntax	The Transfer Syntax of the instance.
CBID	Content Block Identifier	The identifier of the content block being transferred.
CSIZ	Content Size	The size of the original content stored, in bytes.
BSIZ	Object Size	The size of the managed fixed content object (after compression), in bytes.
RSLT	Result Code	The result of the C–STORE operation: SUCS—successful CANC—canceled TOUT—timed-out due to inactivity COMP—presentation contexts not accepted ERRC—lost connection ERFH—failure message sent by remote application entity CTNF—content to be transferred was not found CVRF—content to be transferred failed verification GERR—general error processing content

This audit message means a transfer of content between hosts over a DICOM association completed. The message can be monitored to determine the content sent to particular systems. The “Result Code” field can be used to determine when errors occurred.

DCPS—DICOM C–STORE Start

When a DICOM association initiates a C–STORE operation to transfer content from one host to another, this message is issued.

Table 26: DCPS—DICOM C–STORE Start Fields

Code	Field	Description
ASID	Association Identifier	The unique identifier assigned to the DICOM association.
DIDR	C–STORE Direction	Indicates whether the C–STORE was initiated by the grid node or by a remote host: INBO—initiated by a remote host OUTB—initiated by the node
STUG	Study Instance UID	The Study Identifier of the data being transferred.
SERG	Series Instance UID	The Series Identifier of the data being transferred.
IMGG	SOP Instance UID	The Image Identifier of the data being transferred.
STCL	SOP Class	The SOP Class of the instance.
STTX	Transfer Syntax	The Transfer Syntax of the instance.
CBID	Content Block Identifier	The identifier of the content block being transferred.

This audit message means a transfer of content between hosts over a DICOM association has started. The message can be monitored to determine the content sent to particular systems.

DCSF—DICOM C–STORE Fail

When a an association to perform a requested C–STORE cannot be established, or the information required to establish an association to perform a C–STORE cannot be located, the C–STORE operation fails, and this message is issued.

Table 27: DCSF—DICOM C–STORE Fail Fields

Code	Field	Description
SVIP	Destination Service Port	The destination port for the C–STORE operation. If unknown, this field is omitted from the audit message output.

Table 27: DCSF—DICOM C–STORE Fail Fields (cont.)

Code	Field	Description
DAIP	Destination IP Address	The destination IP address for the C–STORE operation. If unknown, this field is omitted from the audit message output.
RMAE	External Application Entity (AE)	The AE-Title of the destination device. If unknown, this field is omitted from the audit message output.
DIDR	C–STORE Direction	Indicates whether the C–STORE was initiated by the grid node or by a remote host: INBO—initiated by a remote host OUTB—initiated by the node
STUG	Study Instance UID	The Study Identifier of the data being transferred. If unknown, this field is omitted from the audit message output.
SERG	Series Instance UID	The Series Identifier of the data being transferred. If unknown, this field is omitted from the audit message output.
IMGG	SOP Instance UID	The Image Identifier of the data being transferred. If unknown, this field is omitted from the audit message output.
STCL	SOP Class	The SOP Class of the instance. If unknown, this field is omitted from the audit message output.
CBID	Content Block Identifier	The identifier of the content block being transferred.
RSLT	Result Code	Why the C–STORE was unable to complete: CBLK—the CBID associated with the image could not be referenced CBNM—CBID associated with the image did not contain metadata, or had invalid metadata in the CMS ASOF—an association could not be established for the C-STORE request. CSDI—extraction error while processing incoming C-STORE transaction data.

This audit message means a transfer of content between hosts over a DICOM association failed. This can be symptomatic of network problems, or indicate attempts to send data to systems that do not support the image SOP Class.

ETAF—Security Authentication Failed

A connection attempt using Transport Layer Security (TLS) has failed.

Table 28: ETAF—Security Authentication Failed Fields

Code	Field	Description
CNID	Connection Identifier	The unique grid identifier for the TCP/IP connection over which the authentication failed.
RUID	User Identity	A service dependent identifier representing the identity of the remote user.
RSLT	Reason Code	The reason for the failure: SCNI—Secure connection establishment failed. CERM—Certificate was missing. CERT—Certificate was invalid. CERE—Certificate was expired. CERR—Certificate was revoked. CSGN—Certificate signature was invlid. CSGU—Certificate signer was unknown. UCRM—User credentials were missing. UCRI—User credentials were invalid. UCRU—User credentials were disallowed. TOUT—Authentication timed out.

When a connection is established to a secure service that uses TLS, the credentials of the remote entity are verified using the TLS profile and additional logic built into the service. If this authentication fails due to invalid, unexpected, or disallowed certificates or credentials, an audit message is logged. This enables queries for unauthorized access attempts and other security-related connection problems.

The message could result from a remote entity having an incorrect configuration, or from attempts to present invalid or disallowed credentials to the system. This audit message should be monitored to detect attempts to gain unauthorized access to the system.

ETCA—TCP/IP Connection Establish

When a connection to a service running on a node is permitted, this message is generated.

Table 29: ETCA—TCP/IP Connection Establish Fields

Code	Field	Description
SEID	Service Identifier	The unique identifier of the service to which the connection was established.
CNDR	Connection Direction	Indicates whether the connection was opened by the grid node or by a remote host: INBO—connection initiated by a remote host, which connected to the node OUTB—connection initiated by the grid node, which connected to a remote host
SVIP	Destination Service Port	The port the connection was established to.
DAIP	Destination IP Address	The IP address the connection was established to.
SAIP	Source IP Address	The IP address the connection was established from (local IP address).
CNID	Connection Identifier	The unique identifier of the connection.
RSLT	Result Code	Connection status: SUCS—connection successfully established

This audit message means an incoming or outgoing TCP/IP connection was successfully established. This does *not* indicate the corresponding user was permitted to use the service - just that they were not rejected. Typically, each service implements additional authentication mechanisms specific to the service type (DICOM, HTTP etc.).

This message can be used to report on external hosts communicating with the system, and to correlate higher level protocol messages back to the IP address initiating the activity. The “Connection Identifier” field allows correlation of audit messages related to actions performed during a session.

ETCC—TCP/IP Connection Close

When the system on either side of an established connection closes the connection (either normally or abnormally), this message is generated.

Table 30: ETCC—TCP/IP Connection Close Fields

Code	Field	Description
CNID	Connection Identifier	The unique identifier of the connection.
INIE	Initiating Entity	The entity causing the connection to be closed: LOCL—the node closed the connection RMOT—the remote entity closed the connection
RSLT	Result Code	Why the connection was closed: SUCS—connection closed at an expected point CLIN—client (remote side) closed the connection at an expected point LOST—connection closed by the remote entity at an unexpected point TOUT—connection timed-out and was closed

This audit message means a TCP/IP connection was closed. When this message is generated, the corresponding connection ID no longer exists, and the associated TCP/IP connection is no longer established.

This message can be used to detect problems within the system, such as network issues over a WAN, or interoperability problems between systems. The “Connection Identifier” field allows correlation of audit messages related to actions performed during a session.

ETCF—TCP/IP Connection Fail

When an attempt to establish a connection to a remote service fails during establishment, this message is generated.

Table 31: ETCF—TCP/IP Connection Fail Fields

Code	Field	Description
SEID	Service Identifier	The unique identifier of the service to which the connection was attempted.

Table 31: ETCF—TCP/IP Connection Fail Fields (cont.)

Code	Field	Description
CNDR	Connection Direction	Indicates the connection attempt was made to a remote host: OUTB—connection initiated by the grid node, which attempted a connection to a remote host
SVIP	Destination Service Port	The port to which the connection attempt was made.
DAIP	Destination IP Address	The IP address to which the connection attempt was made (remote IP address).
SAIP	Source IP Address	The IP address from which the connection attempt was made (local IP address).
CNID	Connection Identifier	The unique identifier of the attempted connection.
RSLT	Result Code	Why the attempted connection failed: CRFU—outgoing connection refused by remote host UNRE—destination (remote host) unreachable ATHF—TCP/IP connection level authentication failure

This audit message means an outgoing connection attempt failed at the lowest level, due to communication problems - the corresponding service was unable to access the remote host, and the TCP/IP connection was not established.

This message can be used to detect system problems such as configuration errors where content is being pushed to unreachable hosts, or where routing problems result in inaccessibility of hosts.

The message can also be used to report on the hosts to which content was pushed. The “Connection Identifier” field allows correlation of audit messages related to actions performed during a session.

FCRE—File Create

When a new *file* is created on the FSG, this logs the creation.

Table 32: FCRE—File Create Fields

Code	Field	Description
FPTH	File Path	The complete path and name of the file that has been created.

This audit message means a new file entry has been added to the FSG directory tree. The content of the file resides on the local FSG cache, and the process of storing it within the grid has initiated.

Directory creation operations on the FSG do not generate audit messages.

FDEL—File Delete

When an existing file entry in the FSG is deleted, this logs the deletion.

Table 33: FDEL—File Delete Fields

Code	Field	Description
FPTH	File Path	The complete path and name of the file that has been deleted.

This audit message means an existing file entry has been deleted from the FSG directory tree. The content of the file residing within the grid is not affected, however the file becomes inaccessible through the FSG.

Deletion of empty directories on the FSG do not generate audit messages.

Deleting a directory triggers an audit message for *each* enclosed file that is deleted.

FMFY—File Modify

When an existing file entry in the FSG is modified (overwritten), this logs the change.

Table 34: FMFY—File Modify Fields

Code	Field	Description
FPTH	File path	The complete path and name of the file being modified.
UUID	Universal Unique ID	The identifier of the original version of the file within the grid.

The original content of the file being changed is retained within the grid at the UUID provided, but can no longer be accessed through the FSG. The content is available through other direct grid interfaces by referencing the UUID number.

The new content of the file is cached in the local FSG, and the process of storing it within the grid is initiated.

FRNM—File Rename

When an existing *file* entry in the FSG is renamed, this logs the change.

Table 35: FRNM—File Rename Fields

Code	Field	Description
OLDP	Original file path	The complete path and name of the (original) file being renamed.
NEWP	New file path	The complete path and name being assigned to the file.

An existing file entry in the FSG directory tree is changing. The content of the file residing within the grid is not affected, however metadata associating the file path and name is changed.

Renaming a *directory* does not trigger any audit messages. The metadata recorded for any enclosed files remains unchanged, indicating the original ingest location only.

FSTG—File Store to Grid

When new content is stored via the FSG, the content is cached locally by the FSG server and is copied into the grid. When the grid confirms it has stored the copy (and is processing it under its business rules for replication), this message is issued.

Table 36: FSTG—File Store to Grid Fields

Code	Field	Description
FPTH	File path	The complete path and name of the file being stored.
FLTP	File Type	Indicates the type of object storage, as processed by the grid's file type detection.
UUID	Universal Unique ID	The identifier of the file content within the grid.
RSLT	Result Code	The result of the storage operation: SUCS—Successfully stored. FTER—Failed extended type verification (will be re-ingested as a generic object). TOUT—Failed due to timeout. ERRC—Failed due to lost connection. GERR—A general error occurred while storing content.

If a failure is logged, the FSG initiates a new storage attempt. Retries continue until successful.

FSWI—File Swap In

A file has been retrieved from the grid for storage in the FSG local cache. Content still resides in the grid.

Table 37: FSWI—File Swap In Fields

Code	Field	Description
FPTH	File path	The complete path and name of the file added to the FSG local cache.
UUID	Universal Unique ID	The identifier of the file content within the grid.

Table 37: FSWI—File Swap In Fields (cont.)

Code	Field	Description
RSLT	Result Code	The result of the file retrieve operation: SUCS—Successfully retrieved. TOUT—Failed due to timeout. ERRC—Failed due to lost connection. GERR—A general error occurred while retrieving the content.

The original content of the file (along with its associated path and file name metadata) is retained within the grid at the UUID provided.

This message indicates that a file not stored in the FSG local cache has been accessed using the FSG. That access may be for the purpose of modification, in which case the FMFY message should also appear in the audit log.

FSWO—File Swap Out

A file has been purged from the FSG local cache. Content still resides in the grid and can be accessed using the FSG.

Table 38: FSWO—File Swap Out Fields

Code	Field	Description
FPTH	File path	The complete path and name of the file dropped from the FSG local cache.
UUID	Universal Unique ID	The identifier of the file content within the grid.

The original content of the file (along with its associated path and file name metadata) is retained within the grid at the UUID provided. The FSG interface can be used to retrieve the content from the grid.

HCPE—HTTP PUT C–STORE End

An object can be stored into the /DICOM namespace over an established HTTP session by initiating a PUT transaction to process and store the content as a DICOM object in the grid. When DICOM object storage has completed, this message is issued.

Table 39: HCPE—HTTP PUT C–STORE End Fields

Code	Field	Description
HSID	Session Identifier	The unique identifier assigned to the HTTP session established to the node.
STUG	Study Instance UID	The Study Identifier of the data being stored.
SERG	Series Instance UID	The Series Identifier of the data being stored.
IMGG	SOP Instance UID	The Image Identifier of the data being stored.
STCL	SOP Class	The SOP Class of the instance.
STTX	Transfer Syntax	The Transfer Syntax of the instance.
CBID	Content Block Identifier	The identifier of the corresponding content block for the successfully stored content. If the store operation was not successful, this field is set to 0.
UUID	Content UUID	The Universal Unique IDentifier assigned to the successfully stored content. If the UUID was not specified, or the store operation failed, this field is set to the NULL UUID.
CSIZ	Content Size	The size of the original content stored, in bytes.
BSIZ	Object Size	The size of the managed fixed content object (after compression), in bytes.
RSLT	Result Code	The result of the DICOM Store operation: SUCS—successful TOUT—timed-out due to inactivity ERRS—session closed or lost while the C–STORE transaction was being performed CTNF—content to be transferred was not found CVRF—content to be transferred failed verification GERR—general error processing content

This audit message means a transfer of content between hosts over an HTTP session completed. This message is generated prior to, and in addition to, the “HTTP PUT Transaction End” audit message.

HCPS—HTTP PUT C–STORE Start

An object can be stored into the /DICOM namespace over an established HTTP session by initiating a PUT transaction to process and store the content as a DICOM object in the grid. When DICOM object storage has been initiated, this message is issued.

Table 40: HCPS—HTTP PUT C–STORE Start Fields

Code	Field	Description
HSID	Session Identifier	The unique identifier assigned to the HTTP session established to the node.
STUG	Study Instance UID	The Study Identifier of the data being stored.
SERG	Series Instance UID	The Series Identifier of the data being stored.
IMGG	SOP Instance UID	The Image Identifier of the data being stored.
STCL	SOP Class	The SOP Class of the instance.
STTX	Transfer Syntax	The Transfer Syntax of the instance
RSLT	Result Code	Status at the time the C–STORE operation was initiated: SUCS—C–STORE transaction successfully initiated

This audit message means a transfer of content between hosts over an HTTP session has been initiated. This message is generated after, and in addition to, the “HTTP PUT Transaction Start” audit message.

HDEL—HTTP DELETE Transaction

When an HTTP client issues a DELETE transaction, a request is made to remove the specified stored content, and this message is issued.

Table 41: HDEL—HTTP DELETE Transaction Fields

Code	Field	Description
HSID	Session Identifier	The unique identifier assigned to the HTTP session established to the node.
OBNS	Object Namespace	The namespace within which the object to be removed resides.
OBPA	Object Path	The path to the object to be removed.
OBNA	Object Name	The name of the object to be removed.

Table 41: HDEL—HTTP DELETE Transaction Fields (cont.)

Code	Field	Description
UUID	Content UUID	The Universal Unique IDentifier assigned to the content requested for removal.
RSLT	Result Code	Result of the DELETE transaction: SUCS—successful RONL—failed (read-only object) ERRS—session closed or lost while the DELETE transaction was being performed CTNF—content to be deleted not found AUTH—transaction terminated due to authorization failure BRQT—malformed DELETE transaction GERR—general error processing content

This audit message indicates the result of a request to delete content. If the specified content exists, it can be identified via the "Content UUID" field. The "Result Code" field can be used to determine when errors occurred.

HGEE—HTTP GET Transaction End

When an HTTP client completes a GET transaction to transfer content from the HTTP server to the HTTP client, this message is issued.

Table 42: HGEE—HTTP GET Transaction End Fields

Code	Field	Description
HSID	Session Identifier	The unique identifier assigned to the HTTP session established to the node.
OBNS	Object Namespace	The namespace within which the requested object resides.
OBPA	Object Path	The path to the requested object.
OBNA	Object Name	The name of the requested object.
CBID	Content Block Identifier	The unique identifier of the corresponding content block requested. If the CBID is unknown, this field is set to 0.
UUID	Content UUID	The Universal Unique IDentifier corresponding to the requested content. If the UUID is unknown, this field is set to the NULL UUID.

Table 42: HGEE—HTTP GET Transaction End Fields (cont.)

Code	Field	Description
RSLT	Result Code	Result of the GET transaction: SUCS—successful TOUT—timed-out due to inactivity ERRS—session closed or lost while the GET transaction was being performed CTNF—content to be transferred not found or generated (404) error CTRD—content requested resulted in a redirect operation CVRF—content to be transferred failed validation AUTH—transaction terminated due to authorization failure GERR—general error processing content

This audit message means a transfer of content to an HTTP client completed. It can be monitored to determine the content sent to particular systems. The "Result Code" field can be used to determine when errors occurred.

HGES—HTTP GET Transaction Start

When an HTTP client initiates a GET transaction to transfer content from the HTTP server to the HTTP client, this message is issued.

Table 43: HGES—HTTP GET Transaction Start Fields

Code	Field	Description
HSID	Session Identifier	The unique identifier assigned to the HTTP session established to the node.
OBNS	Object Namespace	The namespace within which the requested object resides.
OBPA	Object Path	The path to the requested object.
OBNA	Object Name	The name of the requested object.
CBID	Content Block Identifier	The unique identifier of the corresponding content block requested. If the CBID is unknown, this field is set to 0.
UUID	Content UUID	The Universal Unique IDentifier corresponding to the requested content. If the UUID is unknown, this field is set to the NULL UUID.

Table 43: HGES—HTTP GET Transaction Start Fields (cont.)

Code	Field	Description
RSLT	Result Code	Status at the time the request for the GET transaction was initiated: SUCS—GET transaction successfully initiated BRQT—GET transaction malformed

This audit message means a request for transfer of content to an HTTP client has been initiated. It can be monitored to determine the content sent to particular systems.

HHEA—HTTP HEAD Transaction

When an HTTP client initiates a HEAD transaction to request information about stored content, this message is issued.

Table 44: HHEA—HTTP HEAD Transaction Fields

Code	Field	Description
HSID	Session Identifier	The unique identifier assigned to the HTTP session established to the node.
OBNS	Object Namespace	The namespace within which the requested object resides.
OBPA	Object Path	The path to the requested object.
OBNA	Object Name	The name of the requested object.
CBID	Content Block Identifier	The unique identifier of the corresponding content block about which information is being requested. If the CBID is unknown, this field is set to 0.
UUID	Content UUID	The Universal Unique IDentifier corresponding to the content about which information is being requested. If the UUID is unknown, this field is set to the NULL UUID.

Table 44: HHEA—HTTP HEAD Transaction Fields (cont.)

Code	Field	Description
RSLT	Result Code	Result of the HEAD transaction: SUCS—successful CTNF—specified content was not found, or generated (404) error CTRD—content requested resulted in a redirect operation AUTH—transaction terminated due to authorization failure ERRS—session closed or lost while the HEAD transaction was being performed BRQT—HEAD transaction malformed GERR—general error processing content

This audit message means information about a given piece of content was requested by an HTTP client. It can be monitored to determine the content inspected by clients. The "Result Code" field can be used to determine when errors occurred.

HOPT—HTTP OPTIONS Transaction

When an HTTP client initiates an OPTIONS transaction to discover which HTTP transactions can be performed on a given piece of content, this message is issued.

Table 45: HOPT—HTTP OPTIONS Transaction Fields

Code	Field	Description
HSID	Session Identifier	The unique identifier assigned to the HTTP session established to the node.
OBNS	Object Namespace	The namespace within which the specified object resides.
OBPA	Object Path	The path to the specified object.
OBNA	Object Name	The name of the specified object.

Table 45: HOPT—HTTP OPTIONS Transaction Fields (cont.)

Code	Field	Description
RSLT	Result Code	Result of the OPTIONS transaction: SUCS—successful ERRS—session closed or lost while the OPTIONS transaction was being performed AUTH—transaction terminated due to authorization failure BRQT—OPTIONS transaction malformed GERR—general error processing content

This audit message indicates the result of a request for information about the transactions that can be performed on content. The OPTIONS transaction is typically performed to discover if content can be deleted, created, and so on.

HPOE—HTTP POST Transaction End

When a POST transaction initiated by an HTTP client to query available content completes, this message is issued.

Table 46: HPOE—HTTP POST Transaction End Fields

Code	Field	Description
HSID	Session Identifier	The unique identifier assigned to the HTTP session established to the node.
OBNS	Object Namespace	The namespace within which the query is performed.
RSFD	Results Found	The number of found objects matching the query.
RSLT	Result Code	Result of the POST query operation: SUCS—successful TOUT—timed-out due to inactivity ERRS—session closed or lost while the POST transaction was being performed CMLF—malformed query parameters received from client AUTH—transaction terminated due to authorization failure BRQT—invalid POST query (bad request) GERR—general error processing content

This audit message means an HTTP client has initiated and completed a query for stored content in the specified namespace. It can be monitored to determine the content being queried. The “Result Code” field can be used to determine when errors occurred.

The time between the “HTTP POST Transaction Start” and “HTTP POST Transaction End” audit messages tells you how long particular query operations are taking to complete.

HPOS—HTTP POST Transaction Start

When a POST transaction is initiated by an HTTP client to query available content, this message is issued.

Table 47: HPOS—HTTP POST Transaction Start Fields

Code	Field	Description
HSID	Session Identifier	The unique identifier assigned to the HTTP session established to the node.
OBNS	Object Namespace	The namespace within which the query is performed.
RSLT	Result Code	Status at the time the request for the POST transaction was initiated: SUCS—POST transaction initiated successfully BRQT—malformed POST transaction

This audit message means an HTTP client initiated a query for stored content in the specified namespace. It can be monitored to determine the content being queried.

The time between the “HTTP POST Transaction Start” and “HTTP POST Transaction End” audit messages tells you how long particular query operations are taking to complete.

HPUE—HTTP PUT Transaction End

When an HTTP client completes a PUT transaction to transfer content from the HTTP client to the HTTP server (the node), this message is issued.

Table 48: HPUE—HTTP PUT Transaction End Fields

Code	Field	Description
HSID	Session Identifier	The unique identifier assigned to the HTTP session established to the node.
OBNS	Object Namespace	The namespace within which the stored object was handled.
OBPA	Object Path	The path used to store the object.
OBNA	Object Name	The name of the stored object.
CBID	Content Block Identifier	The identifier of the corresponding content block for the successfully-stored content. If the store operation was not successful, this field is set to 0.
UUID	Content UUID	The Universal Unique IDentifier assigned to the successfully stored content. If the UUID was not specified, or the store operation failed, this field is set to the NULL UUID.
CSIZ	Content Size	The size of the original content stored, in bytes.
BSIZ	Object Size	The size of the managed fixed content object (after compression), in bytes.
RSLT	Result Code	The result of the PUT transaction: SUCS—successful TOUT—timed-out due to inactivity ERRS—session closed or lost while the PUT transaction was being performed CMLF—malformed content received from the client STER—storing the content failed AUTH—transaction terminated due to authorization failure GERR—general error processing content

This audit message means a transfer of content from an HTTP client completed. If content was successfully stored, the CBID and/or UUID fields identify it.

This audit message can be monitored to determine the content sent to particular systems. The "Result Code" field can be used to determine when errors occurred.

HPUS—HTTP PUT Transaction Start

When an HTTP client initiates a PUT transaction to transfer content from the HTTP client to the HTTP server (the node), this message is issued.

Table 49: HPUS—HTTP PUT Transaction Start Fields

Code	Field	Description
HSID	Session Identifier	The unique identifier assigned to the HTTP session established to the node.
OBNS	Object Namespace	The namespace within which the stored object should be handled.
OBPA	Object Path	The path to use when storing the object.
OBNA	Object Name	The name of the object to store.
RSLT	Result Code	The status at the time the request for the PUT transaction was initiated: SUCS—PUT transaction initiated successfully BRQT—malformed PUT transaction

This audit message means a transfer of content from an HTTP client has initiated. It can be monitored to determine the content stored using HTTP.

HTSC—HTTP Session Close

When an HTTP client finishes communicating with a remote host and closes the previously-established HTTP session, this message is issued.

Table 50: HTSC—HTTP Session Close Fields

Code	Field	Description
HSID	Session Identifier	The unique identifier assigned to the HTTP session established to the node.

Table 50: HTSC—HTTP Session Close Fields (cont.)

Code	Field	Description
RSLT	Result Code	Why the session was closed: SUCS—session closed normally, without errors TOUT—timed-out by the node, due to inactivity ERRC—lost the connection over which the session was established ERRT—session terminated due to an error occurring on a transaction AUTH—session terminated due to a failed transaction authorization GERR—a general error occurred, causing the session to close

This audit message means an HTTP client closed a previously-established HTTP session. “HTTP Session Close” always corresponds with a previously-issued “HTTP Session Establish” message.

This message should be monitored to determine if there are any repetitive or excessive problems in attempting to establish a session. This could indicate potential communications or interoperability problems related to HTTP client or server implementations.

HTSE—HTTP Session Establish

When an HTTP client establishes an HTTP session, this message is issued.

Table 51: HTSE—HTTP Session Establish Fields

Code	Field	Description
CNID	Connection Identifier	The unique identifier for the connection over which the HTTP session was established.
HSID	Session Identifier	The unique identifier assigned to the HTTP session established to the node.
RSLT	Result Code	Status at the time the session was established: SUCS—session successfully established

This audit message means a remote host (client) successfully established an HTTP session to the node. It can be used to track which hosts the system is communicating with via the HTTP protocol.

RPSB—Replication Session Begin

When a service begins a replication operation—replicating private structured data to a secondary service—this message is generated.

Table 52: RPSB—Replication Session Begin Fields

Code	Field	Description
RPSI	Replication Session ID	The unique identifier of the replication session being started.
RPPI	Previous Session ID	The identifier of the previous replication session (if one exists); zero otherwise.
RPSE	Replication Source Entity	The node ID of the service that is generating the replication session.
RPDE	Replication Destination Entity	The node ID of the service that is accepting the replication session.
RPSC	Start Sequence Count	The replication sequence count of FSG transactionsat which the session starts or resumes.
RSSS	Session Start Reason	The status of the replication session: NEWS—A new session is being established. CONT—A new session is being resumed. RSUM—A previous session is being resumed.
RSLT	Operation Result	The status of the replication operation: SUCS—The replication session started successfully.

This message indicates a replication session is either starting or being resumed. It identifies the primary (originating) and secondary (accepting) services by their node IDs. *Both* the source and destination services report this message.

RPSE—Replication Session End

When a service completes a replication session, this message is generated.

Table 53: RPSE—Replication Session End Fields

Code	Field	Description
RPSI	Replication Session ID	The unique identifier of the replication session that has ended.

Table 53: RPSE—Replication Session End Fields (cont.)

Code	Field	Description
RPPI	Next Session ID	The identifier of the next replication session (if known). If the next session ID is not known, this value is zero (0).
RPSE	Replication Source Entity	The node ID of the service that is generating the replication session.
RPDE	Replication Destination Entity	The node ID of the service that is accepting the replication session.
RPSC	End Sequence Count	The replication sequence count of FSG transactionsthat would be the next value (in another session).
RSSS	Session End Reason	The completion status of the replication session: SUCS—The replication session was closed successfully. UNEX—The session was closed unexpectedly. PAUS—The session was paused (the FSG was shut down). CKPT—The session was stopped for a checkpoint such as a backup. A new session handles remaining replication.
RSLT	Session Result	The result of the replication session: SUCS—The replication session completed successfully. FAIL—The replication session did not complete successfully.

Matching this message with the corresponding RPSB message can indicate the time it took to perform the replication. This message indicates whether the replication session closed normally. *Both* the source and destination services report this message.

SADD—Security Audit Disable

This message indicates the originating service (node ID) has turned off audit message logging; audit messages are no longer being collected or delivered.

Table 54: SADD—Security Audit Disable Fields

Code	Field	Description
AETM	Enable Method	The method used to disable the audit.

Table 54: SADD—Security Audit Disable Fields (cont.)

Code	Field	Description
AEUN	User Name	The user name that executed the command to disable audit logging.

The message implies that logging was previously enabled but has now been disabled. This is typically only used during bulk ingest to improve system performance. Following the bulk activity, auditing is restored (SADE) and the capability to disable auditing is then permanently blocked.

SADE—Security Audit Enable

This message indicates that the originating service (node ID) has restored audit message logging; audit messages are again being collected and delivered.

Table 55: SADE—Security Audit Enable Fields

Code	Field	Description
AETM	Enable Method	The method used to enable the audit.
AEUN	User Name	The user name that executed the command to enable audit logging.

The message implies that logging was previously disabled (SADD) but has now been restored. This is typically only used during bulk ingest to improve system performance. Following the bulk activity, auditing is restored and the capability to disable auditing is then permanently blocked.

SCMT—Object Store Commit

Grid content is not made available or recognized as being stored until it has been committed - meaning it has been stored persistently. Persistently-stored content has been completely written to disk, and has passed related integrity checks. When a content block is committed to storage, this message is issued.

Table 56: SCMT—Object Store Commit Fields

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block committed to permanent storage.
RSLT	Result Code	Status at the time the object was stored to disk: SUCS - object successfully stored

This message means a given content block has been completely stored and verified, and can now be requested. It can be used to track data flow within the system.

SREM—Object Store Remove

When grid content is removed, it is either downgraded to transient status (“removed”) or completely wiped from the system such that no parts of the content remain (“purged”). If content is downgraded to transient status, it may still be accessed until purged from the system.

When a content block is deleted from permanent storage (either removed or purged), this message is issued.

Table 57: SREM—Object Store Remove Fields

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block deleted from permanent storage.
RSLT	Result Code	How the content was deleted: SUCS - content removed (downgraded to transient status) PURG -content purged from the node INTL - the operation succeeded and the content is no longer accessible - but not erased, as the purge interlock is enabled

This audit message means a given content block has been deleted from a node and can no longer be requested directly. The message can be used to track the flow of deleted content within the system.

SVRF—Object Store Verify Fail

Each time content is read from or written to disk, several verification and integrity checks are performed to ensure data being sent to the requesting user is identical to the data originally ingested into the system. If any of these checks fail, the system automatically removes the corrupt data to prevent it from being retrieved again.

When a content block fails the verification process, this message is issued.

Table 58: SVRF—Object Store Verify Fail Fields

Code	Field	Description
CBID	Content Block Identifier	The unique identifier of the content block which failed verification.
RSLT	Result Code	Verification failure type: CRCF - content CRC checks failed HMAC - content HMAC checks failed EHSH - unexpected encrypted content hash PHSH - unexpected original content hash SEQC - incorrect data sequence on disk PERR - invalid structure of disk file DERR - disk error

The "SVRF - Object Store Verify Fail" audit message should be monitored closely. It means a given content block failed verification checks, which can indicate attempts to tamper with content or impending hardware failures.

SVRU—Object Store Verify Unknown

The Local Distribution Router (LDR) storage component continuously scans all files in the object store to schedule content verification. If it detects a file or directory does not match expected naming conventions, it moves the unexpected file(s) to the "garbage" directory, where they can be automatically or manually removed (depending on LDR configuration).

When an unknown or unexpected file is detected in the object store and moved to the "garbage" directory, this message is issued.

Table 59: SVRU—Object Store Verify Unknown Fields

Code	Field	Description
FPTH	File Path	The full path to the unexpected file's original location.

The "SVRU - Object Store Verify Unknown" audit message should be monitored closely. It means unexpected files were detected in the object store. This situation should be investigated immediately to determine how the files were created, as it can indicate attempts to tamper with content or impending hardware failures.

SYSD—Node Stop

When an HP Medical Archive grid service is stopped gracefully, this message is generated to indicate the shutdown was requested.

Table 60: SYSD—Node Stop Fields

Code	Field	Description
RSLT	Clean Shutdown	The nature of the shutdown: SUCS—System was cleanly shutdown.

The message does not indicate if the host server is being stopped, only the reporting service.

SYSU—Node Start

When an HP Medical Archive grid service is started, this message is generated and indicates if the previous shutdown was clean (commanded) or disorderly (unexpected).

Table 61: SYSU—Node Start Fields

Code	Field	Description
RSLT	Clean Shutdown	The nature of the shutdown: SUCS—System was cleanly shutdown. DSDN—System was not cleanly shutdown.

The message does not indicate if the host server was started, only the reporting service.

This message can be used to:

- Detect discontinuity in the audit trail.
- Determine if a service is failing during operation (as the distributed nature of the grid can mask these failures). The Server Manager restarts a failed service automatically.

TACB—Grid Task Action Begin

When a grid task action begins, this message is generated.

Table 62: TACB—Grid Task Action Begin Fields

Code	Field	Description
TSID	Task ID	The unique identifier of the task used to manage the task over its life cycle.
TTYP	Task Type	The type of task.
TSFC	Task Stage	The current stage of the task.
ACNT	Task Action Node ID	The service node ID being requested to perform the task.
ACTT	Task Action	The action being started.
RSLT	Action Status	Status at the time the task action begins: SUCS—The task stage started successfully.

Matching this message with the corresponding TACE message can indicate the time it took to perform a task action.

TACE—Grid Task Action End

When a grid task action completes, this message is generated.

Table 63: TACE—Grid Task Action End Fields

Code	Field	Description
TSID	Task ID	The unique identifier of the task used to manage the task over its life cycle.
TTYP	Task Type	The type of task.

Table 63: TACE—Grid Task Action End Fields (cont.)

Code	Field	Description
TSFC	Task Stage	The current stage of the task.
ACNT	Task Action Node ID	The service node ID being requested to perform the task.
ACTT	Task Action	The action that is being ended.
RSLT	Action Result	The completion status of the task action: SUCS—completed successfully ABRT—aborted FAIL—failed before completion

Matching this message with the corresponding TACB message can indicate the time it took to perform a task action.

TSGC—Grid Task Stage Change

This message indicates the stage of a grid task has changed; either the task is progressing to the next stage, was aborted, or has failed.

Table 64: TSGC—Grid Task Stage Change Fields

Code	Field	Description
TSID	Task ID	The unique identifier of the task used to manage the task over its life cycle.
TTYP	Task Type	The type of task.
TSDC	Task Stage Description	A text description of the next task stage (starting).
TSFC	Task Stage	The current stage of the task.
RSLT	Task Stage Result	The completion status of the previous task stage: SUCS—completed successfully ABRT—aborted FAIL—failed before completion

All actions within a task stage must complete before the stage can complete.

When a new grid task starts, this message is generated with RSLT = SUCS.

TSTC—Grid Task State Change

When a grid task is added, started, paused, canceled, or completed, this message is issued to audit the change in task state.

Table 65: TSTC—Grid Task State Change Fields

Code	Field	Description
TSID	Task ID	The unique identifier of the task used to manage the task over its life cycle.
TTYP	Task Type	The type of task.
TDSC	Task Description	A text description of the task.
TSRC	Task Source	A text identification of the issuer of the task.
TSTS	Task State	The current state of the task: NEWT—Newly added. PEND—Pending. ACTV—Active (running) PAUS—Paused. ROLA—Aborting; performing a rollback. ROLF—Rollback failed. HIST—Historical (task completed, cancelled, or expired). RMVD—Task is now removed.
RSLT	Task Status	The status of the grid task: SUCS—The task successfully entered the current state. ABRT—The task was aborted (rollback if possible). FAIL—The task has failed (rollback if possible). CANC—The task was cancelled (never started). EXPR—The task expired (never started). IVLD—The task is invalid. AUTH—The task is unauthorized. DUPL—The task is a duplicate.

This message is used to determine what tasks have been added, run, and completed, and the result of the completion.

The Task Status serves to indicate why the task is in the current state. If a task ends abnormally (is aborted or fails) and requires a rollback, the reason is retained in the task status (TCTS). The status of the rollback itself is noted within the state (TSTS). The sequence of messages would indicate either:

- ROLA > HIST if the rollback is successful, or
- ROLA > ROLF > HIST if the rollback failed.

Glossary

ADC Administrative Domain Controller—a unit of the HP Medical Archive software that authenticates grid nodes (certificates) and manages interconnections. It maintains grid topology information.

AE title Application Entity Title—the identifier of a DICOM node communicating with other DICOM AEs.

AMS Audit Management System—a unit of the HP Medical Archive software that monitors and logs all audited system events and transactions.

CBID Content Block Identifier—A number that uniquely identifies a piece of content within the HP Medical Archive system.

CIDR Classless Inter-Domain Routing—a method of routing traffic between IP networks that improves flexibility when dividing ranges of IP Addresses into separate networks. CIDR is defined in RFC 1519. Standard notation for a CIDR address range begins with the network address, padded with zero bits on the right, followed by a slash "/" character and a number representing the length in bits of the subnet mask (*prefix*), thus defining the size of the network. For example:

- 192.168.120.0/24 represents the 256 addresses 192.168.120.0 through 192.168.120.255 inclusive. The "/24" indicates a 24-bit subnet mask, leaving 8 bits (0–255) of subnet address space.
- 192.168.212.0/22 represents the 1024 addresses 192.168.212.0 through 192.168.215.255 inclusive. The left-most 22 bits form the mask, leaving 10 bits (0.0–3.255) of subnet address space.

CIFS Common Internet File System—a file system protocol based on SMB (Server Message Block, developed by Microsoft) intended to complement existing protocols such as HTTP, FTP, and NFS.

CLB Connection Load Balancer—a unit of the HP Medical Archive software that directs incoming DICOM traffic based on factors from an ADC.

CMN Configuration Management Node—a unit of the HP Medical Archive software for performing system-wide reconfiguration and Grid Tasks.

CMS Content Management System—a unit of the HP Medical Archive software managing a distributed database catalog of the grid content (metadata) and data duplication according to business rules to provide Information Lifecycle Management (ILM).

content block ID See “CBID”.

DICOM Digital Imaging and COmmunications in Medicine—a standard developed by ACR-NEMA (an alliance of the American College of Radiology and the National Electrical Manufacturer’s Association) for communications between medical imaging devices.

DR Disaster Recovery.

FCS Fixed Content Storage—A class of stored data where the data, once captured, is rarely changed and must be retained for long periods of time in its original form. Typically this includes images, documents, and other data where alterations would reduce the value of the stored information.

flywheeling A clock is running on its own, without tracking a reference source.

FSG File System Gateway—a unit of the HP Medical Archive software that enables standard network file systems to interface with the grid.

Grid Task A managed sequence of actions that are coordinated across a grid to perform a specific function (such as adding new node certificates). Grid Tasks are typically long-term operations that span many entities within the grid.

HPMA HP Medical Archive—a fixed-content storage system from Hewlett-Packard. The solution is sold under the HP brand and is serviced and supported by the HP services/support organization worldwide. The HPMA Solution is powered by Bycast® StorageGRID™ software.

ILM Information Lifecycle Management—a process of managing data by applying business rules to determine storage accessibility and longevity. Software implementing ILM manages data replication, storage resources, distribution, and retention to meet business and regulatory objectives.

instance A DICOM term for an image. One or more instances for a single patient are collected in a “study”.

LAN Local Area Network—a network of interconnected computers that is restricted to a small area, such as a building or campus. A LAN may be considered a node to the Internet or other wide area network.

latency Time duration for processing a transaction or transmitting a unit of data from end to end. When evaluating system performance, both throughput and latency need to be considered. See also: “throughput”.

LDR Local Distribution Router—a unit of the HP Medical Archive software to manage the storage and transmission of content within the grid.

metadata Data that provides information *about* other data.

namespace A set whose elements are unique names. There is no guarantee that a name in one namespace is not repeated in a different namespace.

NFS Network File System—a protocol (developed by SUN Microsystems) that enables access to network files as if they were on local disks.

NMS Network Management System—a unit of the HP Medical Archive software for alarm monitoring and system administration. It provides a web-based interface for managing and monitoring the HPMA system, as well as viewing and reporting on statistics regarding network, DICOM, storage, and many other related attributes for each of the various services and servers.

object store A configured file system on a disk volume. The configuration includes a specific directory structure and resources initialized at system installation.

PACS Picture Archiving and Communication System—a computerized system of patient records management responsible for short and long term (archival) storage of images. Communication with PACS is via DICOM.

PDF Portable Document Format—a file format (developed by Adobe Systems and based on the postscript language) for exchanging documents between computer systems that may have differing operating systems. It is designed to preserve the appearance of the document regardless of the system used to render it.

release The edition of the complete HP Medical Archive system. Contrast with "version" and "revision".

revision The edition of a document. Contrast with "version" and "release".

Samba A suite of programs that implement the Server Message Block (SMB) protocol. It allows files and printers on the host operating system to be shared with other clients. For example, instead of using telnet to log into a Unix machine to edit a file there, a Windows user might connect a drive in Windows Explorer to a Samba server on the Unix machine and edit the file in a Windows editor.

SQL Structured Query Language—an industry standard interface language for managing relational databases. An SQL database is one that supports the SQL interface.

SSM Service Status Monitor—a unit of the HP Medical Archive software that monitors hardware conditions and reports to the NMS. Every server in the grid runs an instance of the SSM.

study A DICOM term for a collection of images (instances) related to an individual patient or subject.

SVG Scalable Vector Graphic—a format for digital images that can be scaled without loss of resolution.

TCP/IP Transmission Control Protocol / Internet Protocol—a process for encapsulating and transmitting packet data over a network. It includes positive acknowledgement of transmissions.

throughput The amount of data that can be transmitted or the number of transactions that can be processed by a system or subsystem in a given period of time. See also: “latency”.

URI Universal Resource Identifier—A generic set of all names or addresses used to refer to resources that can be served from a computer system. These addresses are represented as short text strings.

URL Universal Resource Location—A URI that can be typed into a browser or other client program in order to retrieve/access an object, such that the client software is able to understand how to perform the requested action. (The client, typically a “browser”, often uses a Domain Name Server (DNS) to resolve a URL into an IP address and URI combination.)

UTC A language-independent international abbreviation, UTC is neither English nor French. It means both “Coordinated Universal Time” and “Temps Universel Coordonné”.

UTC refers to the standard time common to every place in the world. It is derived from International Atomic Time (TAI) by the addition of a whole number of “leap seconds” to synchronize it with Universal Time (UT1). UTC is expressed using a 24-hour clock and uses the Gregorian calendar.

UUID Universal Unique IDentifier—A 128-bit number which is guaranteed to be unique.

version The edition of a service within the HP Medical Archive system. Contrast with “release” and “revision”.

WAN Wide Area Network—a network of interconnected computers that covers a large geographic area such as a country. Contrast with LAN.